

Scaled Sparse Linear Regression

BY TINGNI SUN AND CUN-HUI ZHANG

*Department of Statistics and Biostatistics, Hill Center, Busch Campus, Rutgers University,
Piscataway, New Jersey 08854, U.S.A.*

tingni@stat.rutgers.edu czhang@stat.rutgers.edu

SUMMARY

Scaled sparse linear regression jointly estimates the regression coefficients and noise level in a linear model. It chooses an equilibrium with a sparse regression method by iteratively estimating the noise level via the mean residual square and scaling the penalty in proportion to the estimated noise level. The iterative algorithm costs little beyond the computation of a path or grid of the sparse regression estimator for penalty levels above a proper threshold. For the scaled lasso, the algorithm is a gradient descent in a convex minimization of a penalized joint loss function for the regression coefficients and noise level. Under mild regularity conditions, we prove that the scaled lasso simultaneously yields an estimator for the noise level and an estimated coefficient vector satisfying certain oracle inequalities for prediction, the estimation of the noise level and the regression coefficients. These inequalities provide sufficient conditions for the consistency and asymptotic normality of the noise level estimator, including certain cases where the number of variables is of greater order than the sample size. Parallel results are provided for the least squares estimation after model selection by the scaled lasso. Numerical results demonstrate the superior performance of the proposed methods over an earlier proposal of joint convex minimization.

Some key words: Convex minimization; estimation after model selection; iterative algorithm; linear regression; oracle inequality; penalized least squares; scale invariance; variance estimation.

1. INTRODUCTION

This paper concerns the simultaneous estimation of the regression coefficients and noise level in a high-dimensional linear model. High-dimensional data analysis is a topic of great current interest due to the growth of applications where the number of unknowns far exceeds the number of data points. Among statistical models arising from such applications, linear regression is one of the best understood. Penalization, convex minimization and thresholding methods have been proposed, tested with real and simulated data, and proved to control errors in prediction, estimation and variable selection under various sets of regularity conditions. These methods typically require an appropriate penalty or threshold level. A larger penalty level may lead to a simple model with large bias, while a smaller penalty level may lead to a complex noisy model due to overfitting. Scale-invariance considerations and existing theory suggest that the penalty level should be proportional to the noise level of the regression model. In the absence of knowledge of the latter level, cross-validation is commonly used to determine the former. However, cross-validation is computationally costly and theoretically poorly understood, especially for the purpose of variable selection and the estimation of regression coefficients. The penalty level selected by cross-validation is called the prediction-oracle in Meinshausen & Bühlmann (2006),

who gave an example to show that the prediction-oracle solution does not lead to consistent model selection for the lasso.

Estimation of the noise level in high-dimensional regression is interesting in its own right. Examples include quality control in manufacturing and risk management in finance.

Our study is motivated by Städler et al. (2010) and the comments on that paper by Antoniadis (2010) and Sun & Zhang (2010). Städler et al. (2010) proposed to estimate the regression coefficients and noise level by maximizing their joint log-likelihood with an ℓ_1 penalty on the regression coefficients. Their method has a unique solution due to the joint concavity of the log-likelihood under a certain transformation of the unknown parameters. Sun & Zhang (2010) proved that this penalized joint maximum likelihood estimator may result in a positive bias for the estimation of the noise level and compared it with two alternatives. The first is a one-step bias correction of the penalized joint maximum likelihood estimator. The second is an iterative algorithm that alternates between estimating the noise level via the mean residual square and scaling the penalty level in a predetermined proportion to the estimated noise level in the lasso or minimax concave penalized selection paths. In a simulation experiment Sun & Zhang (2010) demonstrated the superiority of the iterative algorithm, compared with the penalized joint maximum likelihood estimator and its bias correction. However, no theoretical results were given for the iterative algorithm. Antoniadis (2010) commented on the same problem from a different perspective by raising the possibility of adding an ℓ_1 penalty to Huber's concomitant joint loss function. See, for example, section 7.7 of Huber & Ronchetti (2009). Interestingly, the minimizer of this penalized joint convex loss is identical to the equilibrium of the iterative algorithm for the lasso path. Thus, the convergence of the iterative algorithm is guaranteed by the convexity.

In this paper, we study Sun & Zhang (2010)'s iterative algorithm for the joint estimation of regression coefficients and the noise level. For the lasso, this is equivalent to jointly minimizing Huber's concomitant loss function with the ℓ_1 penalty, as Antoniadis (2010) pointed out. For simplicity, we call the equilibrium of this algorithm the scaled version of the penalized regression method, for example the scaled lasso or scaled minimax concave penalized selection, depending on the choice of penalty function. Under mild regularity conditions, we prove oracle inequalities for prediction and the joint estimation of the noise level and regression coefficients for the scaled lasso, that imply the consistency and asymptotic normality of the scaled lasso estimator for the noise level. In addition, we prove parallel oracle inequalities for the least squares estimation of the regression coefficients and noise level after model selection by the scaled lasso. We report numerical results on the performance of scaled lasso and other scaled penalized methods, along with that of the corresponding least squares estimator after model selection. These theoretical and numerical results support the use of the proposed method for high-dimensional regression.

We use the following notation throughout the paper. For a vector $v = (v_1, \dots, v_p)$, $|v|_q = (\sum_j |v_j|^q)^{1/q}$ denotes the ℓ_q norm with the usual extensions $|v|_\infty = \max_j |v_j|$ and $|v|_0 = \#\{j : v_j \neq 0\}$. For design matrices X and subsets A of $\{1, \dots, p\}$, x_j denotes column vectors of X and X_A denotes the matrix composed of columns with indices in A . Moreover, $x_+ = \max(x, 0)$.

2. AN ITERATIVE ALGORITHM

Suppose we observe a design matrix $X = (x_1, \dots, x_p) \in \mathbb{R}^{n \times p}$ and a response vector $y \in \mathbb{R}^n$. For penalty functions $\rho(\cdot)$, consider penalized loss functions of the form

$$L_\lambda(\beta) = \frac{|y - X\beta|_2^2}{2n} + \lambda^2 \sum_{j=1}^p \rho(|\beta_j|/\lambda) \quad (1)$$

where $\beta = (\beta_1, \dots, \beta_p)'$ is a vector of regression coefficients. Let the penalty $\rho(t)$ be standardized to $\dot{\rho}(0+) = 1$, where $\dot{\rho}(t) = (d/dt)\rho(t)$. A vector $\hat{\beta} = (\hat{\beta}_1, \dots, \hat{\beta}_p)'$ is a critical point of the penalized loss (1) if and only if

$$\begin{cases} x'_j(y - X\hat{\beta})/n = \lambda \text{sgn}(\hat{\beta}_j)\dot{\rho}(|\hat{\beta}_j|/\lambda), & \hat{\beta}_j \neq 0, \\ x'_j(y - X\hat{\beta})/n \in \lambda[-1, 1], & \hat{\beta}_j = 0. \end{cases} \quad (2)$$

If the penalized loss (1) is convex in β , then (2) is the Karush–Kuhn–Tucker condition for its minimization.

Given a penalty function $\rho(\cdot)$, one still has to choose a penalty level λ to arrive at a solution of (2). Such a choice may depend on the purpose of estimation, since variable selection may require a larger λ than does prediction. However, scale-invariance considerations and theoretical results suggest using a penalty level proportional to the noise level σ . This motivates a scaled penalized least squares estimator as a numerical equilibrium in the following iterative algorithm:

$$\begin{aligned} \hat{\sigma} &\leftarrow |y - X\hat{\beta}^{\text{old}}|_2 / \{(1-a)n\}^{1/2}, \\ \lambda &\leftarrow \hat{\sigma}\lambda_0, \\ \hat{\beta} &\leftarrow \hat{\beta}^{\text{new}}, \quad L_\lambda(\hat{\beta}^{\text{new}}) \leq L_\lambda(\hat{\beta}^{\text{old}}), \end{aligned} \quad (3)$$

where λ_0 is a prefixed penalty level, not depending on σ , $\hat{\sigma}$ estimates the noise level, and $a \geq 0$ provides an option for a degrees-of-freedom adjustment with $a > 0$. For $p < n$ and $(a, \lambda_0) = (p/n, 0)$, (3) initialized with the least squares estimator $\hat{\beta}^{(\text{lse})}$ is non-iterative and gives $\hat{\sigma}^2 = |y - X\hat{\beta}^{(\text{lse})}|_2^2 / (n - p)$. For large data sets, one may use a few passes of a gradient descent algorithm to compute $\hat{\beta}^{\text{new}}$ from $\hat{\beta}^{\text{old}}$. For $a = 0$, this algorithm was considered in Sun & Zhang (2010). In Sun & Zhang (2010) and the numerical experiments reported in Section 4, $\hat{\beta}^{\text{new}}$ is a solution of (2) for the given λ . We describe this implementation in the following two paragraphs.

The first step of our implementation is the computation of a solution path $\hat{\beta}(\lambda)$ of (2) beginning from $\hat{\beta}(\lambda) = 0$ for $\lambda = |X'y/n|_\infty$. For quadratic spline penalties $\rho(t)$ with m knots, Zhang (2010) developed an algorithm to compute a linear spline path of solutions $\{\lambda^{(t)} \oplus \hat{\beta}^{(t)} : t \geq 0\}$ of (2) to cover the entire range of λ . This extends the least angle regression solution or lasso path (Osborne et al., 2000a,b; Efron et al., 2004) from $m = 1$ and includes the minimax concave penalty for $m = 2$ and the smoothly clipped absolute deviation penalty (Fan & Li, 2001) for $m = 3$. An R package named *plus* is available for computing the solution paths for these penalties.

The second step of our implementation is the iteration (3) along the solution path $\beta(\lambda)$ computed in the first step. That is to use the already computed

$$\hat{\beta}^{\text{new}} = \hat{\beta}(\lambda) \quad (4)$$

in (3). For the scaled lasso, we use $a = 0$ in (3) and $\rho(t) = t$ in (1) and (2). For the scaled minimax concave penalized selection, we use $a = 0$ and the minimax concave penalty $\rho(t) = \int_0^t (1 - x/\gamma)_+ dx$, where $\gamma > 0$ regularizes the maximum concavity of the penalty. When $\gamma = \infty$, it becomes the scaled lasso. The algorithm (3) can be easily implemented once a solution path is computed.

Consider the ℓ_1 penalty. As discussed in the introduction, (3) and (4) form an alternating minimization algorithm for the penalized joint loss function

$$L_{\lambda_0}(\beta, \sigma) = \frac{|y - X\beta|_2^2}{2n\sigma} + \frac{(1-a)\sigma}{2} + \lambda_0|\beta|_1. \quad (5)$$

Antoniadis (2010) suggested this jointly convex loss function as a way of extending Huber's robust regression method to high dimensions. For $a = 0$ and $\lambda = \hat{\sigma}\lambda_0$ with fixed $\hat{\sigma}$, $\hat{\sigma}L_{\lambda_0}(\beta, \hat{\sigma}) = L_{\lambda}(\beta) + \hat{\sigma}^2/2$, so that $\hat{\beta} \leftarrow \hat{\beta}(\lambda)$ in (4) minimizes $L_{\lambda_0}(\beta, \hat{\sigma})$ over β . For fixed $\hat{\beta}$, $\hat{\sigma}^2 \leftarrow |y - X\hat{\beta}|_2^2 / \{(1-a)n\}$ in (3) minimizes $L_{\lambda_0}(\hat{\beta}, \sigma)$ over σ . During the revision of this paper, we learned that She & Owen (2011) have considered penalizing Huber's concomitant loss function for outlier detection in linear regression. We summarize some properties of the algorithm (3) with (4) in the following proposition.

PROPOSITION 1. *Let $\hat{\beta} = \hat{\beta}(\lambda)$ be a solution path of (2) with $\rho(t) = t$. The penalized loss function (5) is jointly convex in (β, σ) and the algorithm (3) with (4) converges to*

$$(\hat{\beta}, \hat{\sigma}) = \arg \min_{\beta, \sigma} L_{\lambda_0}(\beta, \sigma). \quad (6)$$

The resulting estimators $\hat{\beta} = \hat{\beta}(X, y)$ and $\hat{\sigma} = \hat{\sigma}(X, y)$ are scale equivariant in y in the sense that $\hat{\beta}(X, cy) = c\hat{\beta}(X, y)$ and $\hat{\sigma}(X, cy) = |c|\hat{\sigma}(X, y)$. Moreover,

$$\frac{\partial}{\partial \sigma} L_{\lambda_0}\{\hat{\beta}(\sigma\lambda_0), \sigma\} = \frac{1-a}{2} - \frac{|y - X\hat{\beta}(\sigma\lambda_0)|_2^2}{2n\sigma^2}. \quad (7)$$

Since (5) is not strictly convex, the joint estimator may not be unique for some data (X, y) . However, since (5) is strictly convex in σ , $\hat{\sigma}$ is always unique in (6) and the uniqueness of $\hat{\beta}$ follows from that of the lasso estimator $\hat{\beta}(\lambda)$ at $\lambda = \hat{\sigma}\lambda_0$; $\hat{\beta}(\lambda)$ is unique when the second part of (2) is strict in the sense of not hitting $\pm\lambda$ when $\hat{\beta}_j = 0$, which holds almost everywhere in (X, y) for $\lambda > 0$. See, for example, Zhang (2010).

Let $\hat{\sigma}(\lambda) = |y - X\hat{\beta}(\lambda)|_2 / \{(1-a)n\}^{1/2}$. For $\lambda_0 = \{(2/n) \log p\}^{1/2}$, (7) implies that

$$\hat{\sigma} = \hat{\sigma}(\hat{\lambda}), \quad \hat{\lambda} = \min \{\lambda : \hat{\sigma}^2(\lambda) \leq n\lambda^2 / (2 \log p)\}. \quad (8)$$

While the present paper continues our earlier work (Sun & Zhang, 2010) by providing further theoretical and numerical justifications for (3) and (4), the estimator has appeared in different forms. In addition to (5) and (6) of Antoniadis (2010), (8) appeared in Zhang (2010). While this paper was in revision, a reviewer called our attention to Belloni et al. (2011), who focused on studying $\hat{\beta}$ in an equivalent form as square-root lasso. We note that (3) and (4) allow concave penalties and degrees of freedom adjustments as in Zhang (2010).

3. THEORETICAL RESULTS

3.1. Analysis of scaled lasso

Let β^* be a vector of true regression coefficients. An expert with oracular knowledge of β^* would estimate the noise level by the oracle estimator

$$\sigma^* = |y - X\beta^*|_2 / n^{1/2}. \quad (9)$$

Under the Gaussian assumption, this is the maximum likelihood estimator for σ when β^* is known and $n(\sigma^*/\sigma)^2$ follows the χ_n^2 distribution. Due to the scale equivariance of $\hat{\sigma}$ in Proposition 1, it is natural to use σ^* as an estimation target with or without the Gaussian assumption. We derive upper and lower bounds for $\hat{\sigma}/\sigma^* - 1$ and use them to prove the consistency and asymptotic normality of $\hat{\sigma}$. We derive oracle inequalities for the prediction performance and the estimation of β under the ℓ_q loss. Throughout the sequel, $\text{pr}_{\beta, \sigma}$ is the probability measure under which $y - X\beta \sim N(0, \sigma^2 I_n)$. We assume that $|x_j|_2^2 = n$ whenever $\text{pr}_{\beta, \sigma}$ is invoked. The

asymptotic theory here concerns $n \rightarrow \infty$ and allows all parameters and variables to depend on n , including $p \geq n \geq |\beta|_0 \rightarrow \infty$.

We first provide the consistency for the estimation of σ via an oracle inequality for the prediction error of the scaled lasso. In our first theorem, the relative error for the estimation of σ is bounded by a quantity τ_0 related to a prediction error bound $\eta(\lambda, \xi, w, T)$ in (10) below. For $\lambda > 0$, $\xi > 1$, $w \in \mathbb{R}^p$, and $T \subset \{1, \dots, p\}$, define $\delta_{w,T} = 1 - I(w = \beta^*, T = \emptyset)$ and

$$\eta(\lambda, \xi, w, T) = |X\beta^* - Xw|_2^2/n + (1 + \delta_{w,T})2\lambda|w_{T^c}|_1 + \frac{4\xi^2\lambda^2|T|}{(\xi + 1)^2\kappa^2(\xi, T)} \quad (10)$$

where $\kappa(\xi, T)$, the compatibility factor (van de Geer & Bühlmann, 2009), is defined as

$$\kappa(\xi, T) = \min \left\{ \frac{|T|^{1/2}|Xu|_2}{n^{1/2}|u_T|_1} : u \in \mathcal{C}(\xi, T), u \neq 0 \right\} \quad (11)$$

with the cone $\mathcal{C}(\xi, T) = \{u : |u_{T^c}|_1 \leq \xi|u_T|_1\}$. Since the prediction error bound $\eta(\lambda, \xi, w, T)$ is valid for all w and T , τ_0 is related to its minimum over all w and T at the oracle scale σ^* :

$$\tau_0 = \eta_*^{1/2}(\sigma^*\lambda_0, \xi)/\sigma^*, \quad \eta_*(\lambda, \xi) = \inf_{w,T} \eta(\lambda, \xi, w, T). \quad (12)$$

THEOREM 1. *Let $(\hat{\beta}, \hat{\sigma})$ be the scaled lasso estimator in (6), $\beta^* \in \mathbb{R}^p$, σ^* the oracle noise level in (9), $z^* = |X'(y - X\beta^*)/n|_\infty/\sigma^*$ and $\xi > 1$. When $z^* \leq (1 - \tau_0)\lambda_0(\xi - 1)/(\xi + 1)$,*

$$\max \left(1 - \frac{\hat{\sigma}}{\sigma^*}, 1 - \frac{\sigma^*}{\hat{\sigma}} \right) \leq \tau_0, \quad \frac{|X\hat{\beta} - X\beta^*|_2}{n^{1/2}\sigma^*} \leq \frac{1}{\sigma^*}\eta_*^{1/2} \left(\frac{\sigma^*\lambda_0}{1 - \tau_0}, \xi \right) \leq \frac{\tau_0}{1 - \tau_0}. \quad (13)$$

In particular, if $\lambda_0 = A\{(2/n)\log p\}^{1/2}$ with $A > (\xi + 1)/(\xi - 1)$ and $\eta_(\sigma\lambda_0, \xi)/\sigma \rightarrow 0$, then*

$$\text{pr}_{\beta^*, \sigma}(|\hat{\sigma}/\sigma - 1| > \epsilon) \rightarrow 0 \quad (14)$$

for all $\epsilon > 0$.

Theorem 1 extends to the scaled lasso a unification of prediction oracle inequalities for a fixed penalty. With $\lambda = \sigma^*\lambda_0/(1 - \tau_0)_+$, (13) gives $\max\{(\sigma^*\tau_0)^2, |X\hat{\beta} - X\beta^*|_2^2/n\} \leq \eta_*(\lambda, \xi)$, or

$$\max\{(\sigma^*\tau_0)^2, |X\hat{\beta} - X\beta^*|_2^2/n\} \leq \min_w \left\{ |Xw - X\beta^*|_2^2/n + 4\tilde{C}\lambda \sum_{j=1}^p \min(\lambda, |w_j|) \right\} \quad (15)$$

for a $\tilde{C} \geq 1$, if the minimum in (15) is attained at a \tilde{w} with $(1 + 1/\xi)^2\kappa^2(\xi, \tilde{T}) \geq 1/\tilde{C}$, where $\tilde{T} = \{j : |\tilde{w}_j| > \lambda\}$. This asserts that for an arbitrary, possibly non-sparse β^* , the prediction error of the scaled lasso is no greater than that of the best linear predictor Xw with a sparse w for an additional capped- ℓ_1 cost of the order $\lambda \sum_j \min(\lambda, |w_j|)$. A consequence of this prediction error bound for the scaled lasso is the consistency of the corresponding estimator of the noise level in (14). Due to the scale equivariance in Proposition 1, Theorem 1 and the results in the rest of the section are all scale free.

For fixed penalty λ , the upper bound $\eta(\lambda, \xi, w, T)$ has been previously established for different w and T , with possibly different constant factors. Examples include $\eta(\lambda, \xi, \beta^*, \emptyset) = 2\lambda|\beta^*|_1$ (Greenshtein & Ritov, 2004; Greenshtein, 2006), $\eta(\lambda, \xi, \beta^*, S_{\beta^*}) \lesssim \lambda^2|\beta^*|_0$ with $S_w = \{j : w_j \neq 0\}$ (van de Geer & Bühlmann, 2009), and $\min_w \eta(\lambda, \xi, w, S_w) = \min_w \{|X\beta^* - Xw|_2^2/n + O(\lambda^2|w|_0)\}$ (Koltchinskii et al., 2011). In (10), the coefficient for $|Xw - X\beta^*|_2^2/n$ is 1 as in Koltchinskii et al. (2011).

Now we provide sharper convergence rates and the asymptotic normality for the scaled lasso estimation of the noise level σ . This sharper rate $\lambda\mu(\lambda, \xi)/\sigma^2$, essentially taking the square of the order τ_0 in (13), is based on the following ℓ_1 error bound for the estimation of β ,

$$\mu(\lambda, \xi) = (\xi + 1) \min_T \inf_{0 < \nu < 1} \max \left[\frac{|\beta_{T^c}^*|_1}{\nu}, \frac{\lambda|T|/\{2(1-\nu)\}}{\kappa^2\{(\xi+\nu)/(1-\nu), T\}} \right]. \quad (16)$$

This ℓ_1 error bound has the interpretation

$$|\hat{\beta} - \beta^*|_1 \leq \mu(\lambda, \xi) \leq \tilde{C} \sum_{j=1}^p \min(\lambda, |\beta_j^*|), \quad (17)$$

if $\tilde{C} \geq (1 + \xi) \max\{2, 1/\kappa^2(2\xi + 1, \tilde{T})\}$ with $\tilde{T} = \{j : |\beta_j^*| > \lambda\}$. This allows β^* to have many small elements, as in Zhang & Huang (2008), Zhang (2009) and Ye & Zhang (2010). The bound $\mu(\lambda, \xi) \leq (\xi + 1)\lambda|S_{\beta^*}|/\{2\kappa^2(\xi, S_{\beta^*})\}$ improves upon its earlier version in van de Geer & Bühlmann (2009) by a constant factor $4\xi/(\xi + 1) \in (2, 4)$.

THEOREM 2. *Let $\hat{\beta}, \hat{\sigma}, \beta^*, \sigma^*, z^*$ and ξ be as in Theorem 1. Set $\tau_* = \{\lambda_0\mu(\sigma^*\lambda_0, \xi)/\sigma^*\}^{1/2}$. (i) The following inequalities hold when $z^* \leq (1 - \tau_*^2)\lambda_0(\xi - 1)/(\xi + 1)$,*

$$\max(1 - \hat{\sigma}/\sigma^*, 1 - \sigma^*/\hat{\sigma}) \leq \tau_*^2, \quad |\hat{\beta} - \beta^*|_1 \leq \mu(\sigma^*\lambda_0, \xi)/(1 - \tau_*^2). \quad (18)$$

(ii) Let $\lambda_0 \geq \{(2/n)\log(p/\epsilon)\}^{1/2}(\xi + 1)/\{(\xi - 1)(1 - \tau_*^2)\}$. For all $\epsilon > 0$ and $n - 2 > \log(p/\epsilon) \rightarrow \infty$,

$$\text{pr}_{\beta^*, \sigma}\{z^* \leq (1 - \tau_*^2)\lambda_0(\xi - 1)/(\xi + 1)\} \geq 1 - \{1 + o(1)\}\epsilon/\{\pi \log(p/\epsilon)\}^{1/2}.$$

If $\lambda_0 = A\{(2/n)\log p\}^{1/2}$ with $A > (\xi + 1)/(\xi - 1)$ and $\lambda_0\mu(\sigma\lambda_0, \xi)/\sigma \ll n^{-1/2}$, then

$$n^{1/2}(\hat{\sigma}/\sigma - 1) \rightarrow N(0, 1/2) \quad (19)$$

in distribution under $\text{pr}_{\beta^*, \sigma}$.

Since $\sigma^2\tau_*^2 \approx \mu(\lambda, \xi) \leq 2(\xi + 1)\min_T \eta(\lambda, 2\xi + 1, \beta^*, T)$ with $\lambda = \sigma\lambda_0$, the rate τ_*^2 in (18) is essentially the square of that in (13), in view of (12). It follows that the scaled lasso provides a faster convergence rate than does the penalized maximum likelihood estimator for the estimation of the noise level (Städler et al., 2010; Sun & Zhang, 2010). In particular, (18) implies that

$$\max(1 - \hat{\sigma}/\sigma^*, 1 - \sigma^*/\hat{\sigma}) \leq (\xi + 1)\lambda_0^2|S_{\beta^*}|/\{2\kappa^2(\xi, S_{\beta^*})\} \lesssim |\beta^*|_0(\log p)/n \quad (20)$$

with $S_{\beta^*} = \{j : \beta_j^* \neq 0\}$, when $\kappa^2(\xi, S_{\beta^*})$ can be treated as a constant. The bounds in (20) and its general version (18) lead to the asymptotic normality (19) under proper assumptions. Thus, statistical inference about σ is justified with the scaled lasso in certain large- p -smaller- n cases, for example, when $|\beta^*|_0(\log p)/\sqrt{n} \rightarrow 0$ under the compatibility condition (van de Geer & Bühlmann, 2009).

For a fixed penalty level, oracle inequalities for the ℓ_q error of the lasso have been established in Bunea et al. (2007), van de Geer (2008) and van de Geer & Bühlmann (2009) for $q = 1$, Zhang & Huang (2008) and Bickel et al. (2009) for $q \in [1, 2]$, Meinshausen & Yu (2009) for $q = 2$, and Zhang (2009) and Ye & Zhang (2010) for $q \geq 1$. The bounds on $\hat{\sigma}/\sigma^*$ in (18) and (20) allow automatic extensions of these existing ℓ_q oracle inequalities from the lasso with fixed penalty to the scaled lasso. We illustrate this by extending the oracle inequalities of Ye & Zhang (2010) for the lasso and Candes & Tao (2007) for the Dantzig selector in the following corollary. Ye &

Zhang (2010) used the following sign-restricted cone invertibility factor to separate conditions on the error $y - X\beta^*$ and design X in the derivation of error bounds for the lasso:

$$F_q(\xi, S) = \inf \left\{ \frac{|S|^{1/q} |X'Xu|_\infty}{n|u|_q} : u \in \mathcal{C}_-(\xi, S) \right\}, \quad (21)$$

where $\mathcal{C}_-(\xi, S) = \{u : |u_{S^c}|_1 \leq \xi |u_S|_1 \neq 0, u_j x'_j Xu \leq 0, \text{ for all } j \notin S\}$. The quantity (21) can be viewed as a generalized restricted eigenvalue comparing the ℓ_q loss and the dual norm of the ℓ_1 penalty with respect to the inner product for the least squares fit. This gives a direct connection to the Karush–Kuhn–Tucker condition (2). Compared with the restricted eigenvalue (Bickel et al., 2009) and the compatibility factor (11), a main advantage of (21) is to allow all $q \in [1, \infty]$. In addition, (21) yields sharper oracle inequalities (Ye & Zhang, 2010). For $(|A|, |B|, |u|_2) = (\lceil a \rceil, \lceil b \rceil, 1)$ with $A \cap B = \emptyset$, define

$$\delta_a^\pm = \max_{A,u} \left\{ \pm \left(|X_A u|_2 / n^{1/2} - 1 \right) \right\}, \quad \theta_{a,b} = \max_{A,B,u} |X'_A X_B u|_2. \quad (22)$$

The quantities in (22) are used in the uniform uncertainty principle (Candes & Tao, 2007) and the sparse Riesz condition (Zhang & Huang, 2008). We note that $1 - \delta_a^-$ is the minimum eigenvalue of $X'_A X_A / n$ among $|A| \leq a$, $1 + \delta_a^+$ is the corresponding maximum eigenvalue, and $\theta_{a,b}$ is the maximum operator norm of size $a \times b$ off-diagonal sub-blocks of the Gram matrix $X'X/n$.

COROLLARY 1. *Suppose $|\beta_{S^c}^*|_1 = 0$. Then, Theorem 2 holds with $\mu(\lambda, \xi)$ replaced by $\lambda |S| (2\xi) / \{(\xi + 1)F_1(\xi, S)\}$, and for $z^* \leq (1 - \tau_*^2)\lambda_0(\xi - 1)/(\xi + 1)$,*

$$|\hat{\beta} - \beta^*|_q \leq \frac{k^{1/q}(\sigma^* z^* + \hat{\sigma} \lambda_0)}{F_q(\xi, S)} \leq \frac{2\sigma^* \xi \lambda_0 k^{1/q}}{(1 - \tau_*^2)(\xi + 1)F_q(\xi, S)} \quad (23)$$

for all $1 \leq q \leq \infty$, where $k = |S|$. In particular, for $\xi = \sqrt{2}$ and $z^* \leq (1 - \tau_*^2)\lambda_0(\sqrt{2} - 1)^2$,

$$|\hat{\beta} - \beta^*|_2 \leq \frac{(8k)^{1/2} \lambda_0 \sigma^* / (1 - \tau_*^2)}{(\sqrt{2} + 1)F_2(\sqrt{2}, S)} \leq \frac{4k^{1/2} \lambda_0 \sigma^* / (1 - \tau_*^2)}{(\sqrt{2} + 1)(1 - \delta_{1.5k}^- - \theta_{2k, 1.5k})_+}. \quad (24)$$

The proofs of Theorems 1 and 2 are based on a basic inequality

$$\begin{aligned} & |X\hat{\beta}(\lambda) - X\beta^*|_2^2/n + |X\hat{\beta}(\lambda) - Xw|_2^2/n \\ & \leq |Xw - X\beta^*|_2^2/n + 2\lambda\{|w|_1 - |\hat{\beta}(\lambda)|_1\} + 2\sigma^* z^* |w - \hat{\beta}(\lambda)|_1 \end{aligned} \quad (25)$$

as a consequence of the Karush–Kuhn–Tucker conditions (2). The version of (25) with $w = \beta^*$ is well-known (van de Geer & Bühlmann, 2009) and controls $|X\hat{\beta}(\lambda) - X\beta^*|_2^2$ for sparse β^* . When $|X\hat{\beta}(\lambda) - X\beta^*|_2^2 > |Xw - X\beta^*|_2^2$, (25) controls the excess for sparse w by the same argument. The general w is taken in Theorem 1, while $w = \beta^*$ is taken in Theorem 2. In both cases, (25) provides the cone condition in (11) and (21). This is used to derive upper and lower bounds for (7), the derivative of the profile loss function $L_{\lambda_0}(\hat{\beta}(\sigma\lambda_0), \sigma)$ with respect to σ , within a neighborhood of $\sigma/\sigma^* = 1$. The bounds for the minimizer $\hat{\sigma}$ then follow from the joint convexity of the penalized loss (5).

3.2. Estimation after model selection

We have proved that without requiring the knowledge of σ , the scaled lasso enjoys prediction and estimation properties comparable to the best known theoretical results for known σ , and the scaled lasso estimate of σ enjoys consistency and asymptotic normality properties under proper conditions. However, the lasso estimator may have substantial bias (Fan & Peng, 2004; Zhang, 2010), and its bias is significant in our own simulation experiments. Although the smoothly

clipped absolute deviation and minimax concave penalized selectors were introduced to remove the bias of the lasso (Fan & Li, 2001; Zhang, 2010), a theoretical study of their scaled version (3) is beyond the scope of this paper. In this subsection, we present theoretical results for another bias removing method: least squares estimation after model selection.

Given an estimator $\hat{\beta}$ of the coefficient vector β , the least squares estimator of β and the corresponding estimator of the noise level σ in the model selected by $\hat{\beta}$ are

$$\bar{\beta} = \arg \min_{\beta} \left\{ |y - X\beta|_2^2 : \text{supp}(\beta) \subseteq \text{supp}(\hat{\beta}) \right\}, \quad \bar{\sigma} = |y - X\bar{\beta}|_2 / \sqrt{n}, \quad (26)$$

where $\text{supp}(\beta) = \{j : \beta_j \neq 0\}$. Alternatively, we may use $\bar{\sigma} = |y - X\bar{\beta}|_2 / \sqrt{(n - |\hat{\beta}|_0)}$ to estimate the noise level. However, since the effect of this degrees of freedom adjustment is of smaller order than our error bound, we will focus on the simpler (26).

In addition to the compatibility factor $\kappa(\xi, S)$ in (11), we use sparse eigenvalues to study the least squares estimation after the scaled lasso selection. Let $\lambda_{\min}(M)$ be the smallest eigenvalue of a matrix M and $\lambda_{\max}(M)$ the largest. For models $T \subset \{1, \dots, p\}$, define

$$\kappa_{-}(m, T) = \min_{B \supset T, |B \setminus T| \leq m} \lambda_{\min}(X_B' X_B / n), \quad \kappa_{+}(m, T) = \min_{B \cap T = \emptyset, |B| \leq m} \lambda_{\min}(X_B' X_B / n),$$

as the sparse lower eigenvalue of the Gram matrix for models containing T and the sparse upper eigenvalue for models disjoint with T . Let $S = \text{supp}(\beta^*)$ and $\hat{S} = \text{supp}(\hat{\beta})$. The following theorem provides prediction and estimation error bounds for (26) after the scaled lasso selection, along with an upper bound for the false positive $|\hat{S} \setminus S|$, a key element in our study.

THEOREM 3. *Let $(\hat{\beta}, \hat{\sigma})$ be the scaled lasso estimator in (6) and $(\bar{\beta}, \bar{\sigma})$ the least squares estimator (26) in model \hat{S} . Let $\beta^*, \sigma^*, z^*, \xi$ and τ_* be as in Theorem 2 and m be an integer satisfying $|S|\xi^2 / \kappa^2(\xi, S) < m / \kappa_{+}(m, S)$. If $z^* \leq (1 - \tau_*^2)\lambda_0(\xi - 1) / (\xi + 1)$, then*

$$|\hat{S} \setminus S| < m, \quad \hat{\sigma}^2 - \{\sigma_{m-1, S}^* + \sqrt{\eta_*(\hat{\lambda}, \xi)}\}^2 \leq \bar{\sigma}^2 \leq \hat{\sigma}^2, \quad (27)$$

with $\hat{\lambda} = \hat{\sigma}\lambda_0 \leq \sigma^*\lambda_0 / (1 - \tau_*^2)$ and $\sigma_{m, S}^* = \max_{B \supseteq S, |B \setminus S| = m} |(y - X\beta^*)_B|_2 / \sqrt{n}$, and

$$\kappa_{-}(m - 1, S) |\bar{\beta} - \beta^*|_2^2 \leq |X\bar{\beta} - X\beta^*|_2^2 / n \leq \{\sigma_{m-1, S}^* + 2\sqrt{\eta_*(\hat{\lambda}, \xi)}\}^2. \quad (28)$$

Moreover, in addition to the probability bound for $z^* \leq (1 - \tau_*^2)\lambda_0(\xi - 1) / (\xi + 1)$ in Theorem 2 (ii), for all integers $1 \leq m \leq p$,

$$\text{pr}_{\beta^*, \sigma} [\sigma_{m, S}^* (\sqrt{n}) / \sigma \geq \sqrt{(m + |S|)} + \sqrt{\{(2m) \log(ep/m) + 2 \log(1/\epsilon)\}}] \leq \frac{\epsilon/m}{\sqrt{(2\pi)}}. \quad (29)$$

For Gaussian design matrices, the sparse eigenvalues $\kappa_{-}(m, \emptyset)$ and $\kappa_{+}(m, \emptyset)$ can be treated as constants when $m(\log p)/n$ is small and the eigenvalues of the expected Gram matrix are uniformly bounded away from zero and infinity (Zhang & Huang, 2008). Since $\kappa_{-}(m, S) \geq \kappa_{-}(m + |S|, \emptyset)$ and $\kappa_{+}(m, S) \leq \kappa_{+}(m, \emptyset)$, they can be treated as constants in the same sense in Theorem 3. Thus, for sufficiently small $|S|(\log p)/n$, we may take an m of the same order as $|S|$. In this case, the difference between $(\bar{\beta}, \bar{\sigma})$ and the scaled lasso estimator $(\hat{\beta}, \hat{\sigma})$ is of no greater order than the difference between $(\bar{\beta}, \hat{\sigma})$ and the estimation target (β^*, σ^*) . Consequently,

$$|\bar{\sigma}/\sigma - 1| + |\hat{\beta} - \beta^*|_2^2 + |X\bar{\beta} - X\beta^*|_2^2 / n = O_P(1) |S|(\log p)/n.$$

As we have mentioned earlier, the key element in our analysis of (26) is the bound $|\hat{S} \setminus S| < m$ in (27). Since this is a weaker assertion than variable consistency $\hat{S} = S$, the conditions of

Table 1: Performance of five methods in Example 1 at penalty levels $\lambda_j = \{2^{j-1}(\log p)/n\}^{1/2}$ ($j = 1, 2, 3$), across 100 replications, in terms of the bias ($\times 10$) and standard error ($\times 10$) of $\hat{\sigma}/\sigma$ for the selector and $\bar{\sigma}/\sigma$ for the least squares estimator after model selection, the average model size, and the relative frequency of sure screening, along with the simulation results in Fan et al. (2012)

Method		$\hat{\sigma}/\sigma$	$r_0 = 0$		AMS	SSP	$\hat{\sigma}/\sigma$	$r_0 = 0.5$		AMS	SSP
			$\bar{\sigma}/\sigma$					$\bar{\sigma}/\sigma$			
PMLE	λ_1	1.6 \pm 0.6	-1.1 \pm 0.7	7.6	1.0	1.5 \pm 0.6	-1.0 \pm 0.7	9.7	1.0		
	λ_2	2.5 \pm 0.6	-0.1 \pm 0.6	3.0	1.0	2.5 \pm 0.6	-0.3 \pm 0.6	5.2	1.0		
	λ_3	3.6 \pm 0.7	1.2 \pm 1.1	1.8	0.2	3.8 \pm 0.6	-0.2 \pm 0.6	3.7	1.0		
BC	λ_1	0.5 \pm 0.6	-1.8 \pm 0.7	11.9	1.0	0.0 \pm 0.6	-1.9 \pm 0.7	15.5	1.0		
	λ_2	1.6 \pm 0.6	-0.1 \pm 0.6	3.1	1.0	0.7 \pm 0.6	-0.3 \pm 0.6	6.1	1.0		
	λ_3	3.3 \pm 0.7	1.1 \pm 1.1	1.9	0.3	1.9 \pm 0.7	-0.2 \pm 0.6	4.2	1.0		
Scaled lasso	λ_1	0.0 \pm 0.6	-2.1 \pm 0.8	14.6	1.0	-0.5 \pm 0.6	-2.3 \pm 0.7	18.6	1.0		
	λ_2	1.3 \pm 0.7	-0.2 \pm 0.6	3.1	1.0	0.4 \pm 0.6	-0.3 \pm 0.6	6.2	1.0		
	λ_3	3.1 \pm 0.7	1.0 \pm 1.1	1.9	0.3	1.2 \pm 0.7	-0.2 \pm 0.6	4.4	1.0		
Scaled MCP	λ_1	-1.2 \pm 0.8	-2.4 \pm 0.8	14.1	1.0	-0.7 \pm 0.6	-2.2 \pm 0.8	13.8	1.0		
	λ_2	-0.1 \pm 0.6	-0.1 \pm 0.6	3.1	1.0	0.1 \pm 0.6	-0.2 \pm 0.6	3.2	1.0		
	λ_3	1.5 \pm 1.3	0.6 \pm 1.1	2.4	0.6	0.6 \pm 0.7	-0.1 \pm 0.6	3.0	1.0		
Scaled SCAD	λ_1	-0.6 \pm 0.6	-2.2 \pm 0.7	14.0	1.0	-0.4 \pm 0.6	-2.2 \pm 0.8	13.9	1.0		
	λ_2	0.8 \pm 1.0	-0.1 \pm 0.6	3.1	1.0	0.4 \pm 0.6	-0.3 \pm 0.6	3.8	1.0		
	λ_3	3.1 \pm 0.7	0.9 \pm 1.1	2.0	0.3	1.2 \pm 0.7	-0.2 \pm 0.6	3.8	1.0		
N-LASSO		-5.3 \pm 2.0		36.6	1.0	-4.6 \pm 2.0		29.6	1.0		
RCV-SIS		0.2 \pm 1.4		50.0	0.9	-0.1 \pm 1.4		50.0	1.0		
RCV-ISIS		0.5 \pm 1.7		30.9	0.7	0.2 \pm 1.2		29.0	0.8		
RCV-LASSO		0 \pm 1.3		31.1	0.9	-0.3 \pm 1.1		26.5	1.0		
P-SCAD		-1.4 \pm 1.1		30.0	1.0	-1.2 \pm 1.7		29.9	1.0		
CV-SCAD		0.7 \pm 1.2		30.0	1.0	0.9 \pm 1.3		29.9	1.0		
P-LASSO		-0.8 \pm 2.1		36.5	1.0	-0.9 \pm 1.5		29.6	1.0		
CV-LASSO		1.4 \pm 1.1		36.5	1.0	0.8 \pm 1.0		29.6	1.0		

PMLE, ℓ_1 penalized maximum likelihood estimator; BC, bias-corrected estimator; MCP, minimax concave penalty; SCAD, smoothly clipped absolute deviation penalty; N, naive; RCV, refitted cross-validation; SIS, sure independent screening; ISIS, iterative SIS; P, plug-in method with degrees-of-freedom correction; CV, cross-validation; AMS, average model size; SSP, relative frequency of sure screening

Theorem 3 on the design matrix is of a weaker form than the irrepresentability condition for variable selection consistency (Meinshausen & Bühlmann, 2006; Zhao & Yu, 2006). In Zhang & Huang (2008) and Zhang (2010), upper bounds for the false positive were obtained under a sparse Riesz condition on $\kappa_-(m, \emptyset)$ and $\kappa_+(m, \emptyset)$.

4. NUMERICAL RESULTS

4.1. Simulation study

In this section, we present some simulation results to compare five methods: the scaled penalized methods with the ℓ_1 penalty, the minimax concave penalty and the smoothly clipped absolute deviation penalty, the ℓ_1 penalized maximum likelihood estimator (Städler et al., 2010), and its bias correction (Sun & Zhang, 2010). The least squares estimator after model selection by these five methods is also studied. The penalized maximum likelihood estimator is

$$(\hat{\beta}^{(\text{pmle})}, \hat{\sigma}^{(\text{pmle})}) = \arg \max_{\beta, \sigma} \left(-\frac{|y - X\beta|_2^2}{2\sigma^2 n} - \log \sigma - \lambda_0 \frac{|\beta|_1}{\sigma} \right),$$

or equivalently the limit of the iteration $\hat{\sigma} \leftarrow \{y'(y - X\hat{\beta})/n\}^{1/2}$ and $\hat{\beta} \leftarrow \hat{\beta}(\hat{\sigma}\lambda_0)$. The bias-corrected estimator is one iteration of (3) with (4) from $(\hat{\beta}^{(\text{pmle})}, \hat{\sigma}^{(\text{pmle})})$ with $a = 0$,

$$\hat{\sigma}^{(\text{bc})} = |y - \hat{\beta}(\hat{\sigma}^{(\text{pmle})}\lambda_0)|_2/n^{1/2}, \quad \hat{\beta}^{(\text{bc})} = \hat{\beta}(\hat{\sigma}^{(\text{bc})}\lambda_0).$$

Two simulation examples are considered.

Example 1. We compare the five estimators at three penalty levels $\lambda_j = \sqrt{\{2^{j-1}(\log p)/n\}}$, $j = 1, 2, 3$. The experiment has the setting of Example 2 in Fan et al. (2012), with the smallest signal, $b = 1/\sqrt{3}$. We provide their description of the simulation setting in our notation as follows: X has independent and identically distributed Gaussian rows with marginal distribution $N(0, 1)$, $\text{corr}(x_i, x_j) = r_0$ for $1 \leq i < j \leq 50$ and $\text{corr}(x_i, x_j) = 0$ otherwise, $(n, p) = (200, 2000)$, nonzero coefficients $\beta_j = 1/\sqrt{3}$ for $j \in S = \{1, 2, 3\}$, and $y - X\beta \sim N(0, \sigma^2 I)$ with $\sigma = 1$. Two configurations are considered: independent columns x_j with $r_0 = 0$ and correlated first 50 columns x_j with $r_0 = 0.5$. We set $\gamma = 2/(1 - \max |x'_k x_j|/n)$ for the concave penalties.

The top section of Table 1 presents our simulation results, while the bottom section includes the simulation results of Fan et al. (2012) for several joint estimators of (β, σ) using cross-validation, without repeating their experiment. In addition to the bias and the standard error of the ratios $\hat{\sigma}/\sigma$ for the five original estimators and $\bar{\sigma}/\sigma$ for the least squares estimation after model selection, we report the average model size $|\hat{S}|$ and the relative frequency of sure screening, $\hat{S} \supseteq S$, as in Fan et al. (2012), where $\hat{S} = \text{supp}(\hat{\beta})$ is the selected model.

Without post processing, the scaled minimax concave penalized selector with the universal penalty level $\lambda_2 = \sqrt{\{(2/n) \log p\}}$ clearly outperforms other procedures in this example. However, the results of the least squares estimation after model selection at penalty level λ_2 are nearly identical to the top performer for all five methods. In view of the results in average model size and sure screening proportion, the success of post processing at λ_2 is clearly due to the success of model selection. The five methods select too few variables at the larger penalty level λ_3 and too many at the smaller λ_1 , both leading to substantial bias in the estimation of σ for $r_0 = 0$. For $r_0 = 0.5$, selecting a slightly smaller model does not harm so much since a substantial portion of the effect of the missing variables is explained by the selected variables correlated to them. The minimax concave penalized selector is nearly unbiased in this example, so that it does not need post processing. Cross-validation methods select about 30 variables when the true model size is 3. This over selection is probably the reason for the large bias for most cross-validation methods and large standard error for all of them.

Example 2. This experiment has the same setting as in the simulation study in Sun & Zhang (2010), where the scaled lasso and the scaled minimax concave penalized selection are called the

Table 2: Performance of five methods in Example 2 at penalty levels $\lambda_j = \{2^{j-1}(\log p)/n\}^{1/2}$ ($j = 1, 2, 3$), across 100 replications, in terms of the bias ($\times 10$) and standard error ($\times 10$) of $\hat{\sigma}/\sigma$ and $\bar{\sigma}/\sigma$, the false positive, and the false negative

Method		$\hat{\sigma}/\sigma$	$r_0 = 0.1$		FP	FN	$\hat{\sigma}/\sigma$	$r_0 = 0.9$		FP	FN
			$\bar{\sigma}/\sigma$					$\bar{\sigma}/\sigma$			
PMLE	λ_1	5.5 \pm 0.3	0.2 \pm 0.3	3	11	2.4 \pm 0.3	-0.4 \pm 0.3	4	12		
	λ_2	7.7 \pm 0.4	2.1 \pm 0.6	0	19	3.7 \pm 0.3	-0.3 \pm 0.3	1	15		
	λ_3	9.5 \pm 0.4	6.3 \pm 1.1	0	30	5.7 \pm 0.3	-0.1 \pm 0.3	0	20		
BC	λ_1	3.2 \pm 0.3	-0.3 \pm 0.4	8	9	0.3 \pm 0.3	-0.7 \pm 0.3	9	11		
	λ_2	6.1 \pm 0.5	1.6 \pm 0.5	0	17	1.2 \pm 0.3	-0.3 \pm 0.3	2	13		
	λ_3	9.1 \pm 0.5	6.1 \pm 1.1	0	29	3.1 \pm 0.4	-0.1 \pm 0.3	0	18		
Scaled lasso	λ_1	1.9 \pm 0.4	-0.8 \pm 0.4	14	7	0.0 \pm 0.3	-0.8 \pm 0.3	11	11		
	λ_2	5.0 \pm 0.5	1.3 \pm 0.5	0	16	0.6 \pm 0.3	-0.3 \pm 0.3	2	13		
	λ_3	8.9 \pm 0.6	6.0 \pm 1.2	0	29	1.8 \pm 0.4	-0.2 \pm 0.3	0	16		
Scaled MCP	λ_1	-0.2 \pm 0.4	-1.1 \pm 0.4	14	7	0.0 \pm 0.3	-0.7 \pm 0.3	10	19		
	λ_2	1.8 \pm 0.5	0.7 \pm 0.5	0	13	0.6 \pm 0.3	-0.2 \pm 0.3	1	20		
	λ_3	7.8 \pm 1.0	5.6 \pm 1.3	0	28	1.7 \pm 0.4	0.0 \pm 0.3	0	22		
Scaled SCAD	λ_1	0.6 \pm 0.4	-1.0 \pm 0.4	14	7	0.0 \pm 0.3	-0.8 \pm 0.3	10	15		
	λ_2	4.7 \pm 0.6	1.3 \pm 0.5	0	16	0.6 \pm 0.3	-0.3 \pm 0.3	2	15		
	λ_3	8.9 \pm 0.6	6.0 \pm 1.2	0	29	1.8 \pm 0.4	-0.2 \pm 0.3	0	17		

PMLE, ℓ_1 penalized maximum likelihood estimator; BC, bias-corrected estimator; MCP, minimax concave penalty; SCAD, smoothly clipped absolute deviation penalty; FP, false positive; FN, false negative

naive estimators. We provide the description of the simulation settings in Sun & Zhang (2010) in our notation as follows: $(n, p) = (600, 3000)$, the x_j are normalized columns from a Gaussian random matrix with independent and identically distributed rows and correlation $r_0^{|k-j|}$ between the j -th and k -th entries within each row, $\gamma = 2/(1 - \max |x'_k x_j|/n)$ for the minimax concave penalty and smoothly clipped absolute deviation penalty, the nonzero β_j^* are composed of five blocks of $\beta_*(1, 2, 3, 4, 3, 2, 1)'$ centered at random multiples j_1, \dots, j_5 of 25, β_* sets $|X\beta_*|_2^2 = 3n$, and $y - X\beta^*$ is a vector of independent and identically distributed $N(0, 1)$ variables. Thus, the true noise level is $\sigma = 1$. We set $r_0 = 0.1$ for low correlation between design vectors and $r_0 = 0.9$ for high correlation.

We summarize the simulation results in Table 2, which provides the bias and standard error of the ratios $\hat{\sigma}/\sigma$ for the selector and $\bar{\sigma}/\sigma$ for the least squares estimator after model selection, the false positive $|\hat{S} \setminus S|$, and the false negative $|S \setminus \hat{S}|$. Without post processing, the scaled lasso outperforms the penalized maximum likelihood estimator and its bias correction, which are also based on the lasso path. However, the scaled lasso estimate of σ is still biased, and the level of bias is comparable with the order of the error bound $(|S|/n) \log p = 0.47$ in (20). This and the failure in sure screening by any method reflect the difficulty of this example, where $|S| = 35$ is not small and the signal is weak, with average $\beta_* = 0.11$ and 0.05 respectively for $r_0 = 0.1$ and $r_0 = 0.9$. From this perspective, the scaled minimax concave penalized

Table 3: Estimated coefficients ($\times 10^3$) of selected probe sets by four methods in the real data example: the lasso with cross-validation, the lasso with adjusted cross-validation, and the scaled lasso and minimax concave penalized selection at $\lambda_0 = \lambda_2 = \{2(\log p)/n\}^{1/2}$

Probe ID	C-V lasso		C-V lasso/LSE		Scaled lasso		Scaled MCP	
#cov	200	3000	200	3000	200	3000	200	3000
1369353.at	-9.12	-7.13*	-7.09	-2.79*	-7.3	-4.03*		
1370052.at [△]		3.65						
1370429.at	-3.22		-8.94*	-11.06	-8.78*	-9.36	-16.37*	
1371242.at	-6.66							
1374106.at	8.88*	10.58*	7.33*	6.14*	7.45*	7.01*	8.47*	10.02*
1374131.at	4.07	0.80						
1375585.at [△]						0.58		
1384204.at			0.70		0.70			
1387060.at [△]		3.50*						
1388538.at [△]		1.42						
1389584.at	17.16*	25.39*	20.07*	19.61*	19.97*	21.18*	45.75*	50.49*
1393979.at	-1.81		-0.22		-0.4			
1379079.at [△]		-1.43*						
1379495.at		4.84	1.73		1.71	1.00		
1379971.at	13.56*	13.1	11.19*	8.81	11.25*	9.52		
1380033.at	8.69		2.76		2.97		6.75*	
1380070.at [△]		0.19						
1381787.at	-2.05		-2.01		-2.11			
1382452.at [△]		12.93				1.63		12.91*
1382835.at	12.64	5.79	3.73		4.15			
1383110.at	9.03*	19.99	15.10*	16.43	14.97*	16.69	15.80*	23.01*
1383522.at	3.03*		*		*			
1383673.at	5.54	6.12*	6.07	6.15*	6.08	6.47*		
1383749.at	-13.86	-10.85*	-10.84	-6.7*	-11.02	-8.07*	-2.74 *	-1.11*
1383996.at	25.01*	17.82*	18.61*	14.30*	18.88*	15.52*	25.07*	19.19*
1385687.at [△]		-0.99						
1386683.at				4.60*		2.90*		
1390788.a.at	0.92							
1392692.at [△]		1.74						
1393382.at	2.43							
1393684.at	1.59							
1395076.at [△]						0.23		
1397489.at [△]		3.33						
Model size	19	20	15	10	15	14	7	6
$\hat{\lambda} = \hat{\sigma}\lambda_0$	0.0103	0.0163	0.025	0.035	0.0243	0.0315	0.0244	0.0304

C-V lasso/LSE, the lasso estimator with the adjusted cross-validation; MCP, minimax concave penalty; #cov, the number of covariates considered; [△], probes not in the smaller set of 200 probes; *, covariates selected by stability selection

selection, designed to reduce the bias of the lasso, performs quite well at the universal penalty level $\lambda_2 = \sqrt{\{(2/n) \log p\}}$, especially with post processing. The least squares estimation after model selection reduces the bias substantially in all cases, even without successful model selection. This example seems to suggest the possibility of improving the performance of the scaled estimators at a penalty level λ smaller than the universal penalty level λ_2 , a simple upper bound for $|X'(y - X\beta)/(\sigma^*n)|_\infty$ under $\text{pr}_{\beta,\sigma}$. However, consistent variable selection requires




Fig. 1: Mean squared prediction error against penalty level: solid curve, testing error of the lasso for fixed λ ; dotted curve, training error of the lasso for fixed λ ; dot-dashed line, testing error of the scaled lasso with fixed $\lambda_0 = \sqrt{\{(2/n) \log p\}}$, or equivalently the lasso at penalty level $\hat{\sigma} \lambda_0$.

$\lambda \geq |X'(y - X\beta^*)/(\sigma^*n)|_\infty$ as Example 1 demonstrates. Since the scaled lasso estimator $\hat{\sigma}$ is an increasing function of the penalty level by Proposition 1, it is always possible to reduce the bias of $\hat{\sigma}$ to zero by taking a specific λ for each specific example. However, the two examples in our simulation experiment demonstrate the difficulty of picking such a penalty level consistently.

4.2. Real data example

We study a data set containing 18976 probes for 120 rats, which is reported in Scheetz et al. (2006). Our goal is to find probes that are related to that of gene TRIM32, which has been found to cause Bardet-Biedl syndrome, a genetically heterogeneous disease of multiple organ systems including the retina. We consider linear regression with the probe from TRIM32, 1389163_at, as the response variable. As in Huang et al. (2008), we focus on 3000 probes with the largest variances among the 18975 covariates and consider two approaches. The first approach is to regress on these $p = 3000$ probes. The second approach is to regress on the 200 probes among the 3000 with the largest marginal correlation coefficients with TRIM32. For the cross-validation lasso, we randomly partition the data 1000 times, each with a training set of size 80 and a validation set of size 40. For each partition, the penalty level λ is selected by minimizing the prediction mean squared error in the validation set. Then we compute the lasso estimator with all 120 observations at the penalty level equal to the median of the selected penalty levels with the 1000 random partitions. Since cross-validation tends to choose a larger model, we also consider an adjusted version using the cross-validated error of the least squares estimator after the lasso selection. For the minimax concave penalty, we set $\gamma = 2/(1 - \sigma_{0.95}) = 6.37$, where $\sigma_{0.95}$ is the 95% quantile of $|x'_k x_j|/n$.

Table 3 shows the probe sets identified by four methods: the cross-validation lasso, its adjusted version, the scaled lasso at a universal penalty level $\lambda_2 = \{2(\log p)/n\}^{1/2}$, and the minimax concave penalized selection at the same penalty level. We apply stability selection (Meinshausen & Bühlmann, 2010) to check the reliability of selection. Let W_1, \dots, W_p be independent vari-

Table 4: Prediction performance of eight methods in the real data example at penalty levels $\lambda_0 = \lambda_j = \{2^{j-1}(\log p)/n\}^{1/2} (j = 1, 2, 3)$, in terms of the prediction mean squared error ($\times 10^2$), estimated model size, and correlation coefficient ($\times 10^2$) between fitted and observed responses

Method		#cov = 200			#cov = 3000		
		P-MSE	$ \hat{\beta} _0$	corr	P-MSE	$ \hat{\beta} _0$	corr
PMLE	λ_1	0.94	12	67.1	0.97	12	63.5
	λ_2	0.97	9	63.5	1.04	7	59.8
	λ_3	1.09	6	57.6	1.23	3	52.2
BC	λ_1	0.93	13	68.2	0.96	15	64.6
	λ_2	0.95	10	64.7	1.01	9	60.9
	λ_3	1.04	7	59.4	1.17	4	53.1
Scaled lasso	λ_1	0.93	13	68.4	0.96	17	64.3
	λ_2	0.94	10	65.2	0.98	10	61.7
	λ_3	1.02	7	60.8	1.13	5	53.9
Scaled MCP	λ_1	1.03	6	66.4	1.08	8	62.3
	λ_2	1.03	5	63.4	1.06	5	60.0
	λ_3	1.12	3	59.1	1.18	2	54.9
Scaled SCAD	λ_1	1.00	11	68.9	1.01	14	65.1
	λ_2	0.95	10	68.8	0.98	10	65.9
	λ_3	1.01	8	65.0	1.09	5	59.7
C-V lasso		0.94	15	69.0	0.99	25	63.8
C-V lasso/LSE1		0.97	11	64.8	0.98	12	62.5
C-V lasso/LSE2		0.97	11	66.8	1.09	12	62.6

PMLE, ℓ_1 penalized maximum likelihood estimator; BC, bias-corrected PMLE; MCP, minimax concave penalty; SCAD, smoothly clipped absolute deviation penalty; C-V lasso/LSE1, the lasso with adjusted cross-validation; C-V lasso/LSE2, the least squares estimator after the lasso selection with adjusted cross-validation, #cov, the number of covariates considered; corr, the correlation coefficient between fitted and observed responses; P-MSE, prediction mean squared error

ables with $P(W = 0.2) = P(W = 1) = 1/2$ and

$$\hat{\beta}^W = \arg \min_b \frac{|y - Xb|_2^2}{2n} + \hat{\lambda} \sum_{j=1}^p |\beta_j|/W_j,$$

where $\hat{\lambda}$ is the penalty level chosen by individual methods. Stability selection selects variables with nonzero estimated $\hat{\beta}_j^W$ over 50 times in 100 replications. We observe that the scaled minimax concave penalized selector produces most sparse and most stable selection, followed by the adjusted cross-validation, the scaled lasso and then the plain cross-validation. The selection results are consistent among the four methods in the sense that the selected models are almost nested. Since the model size is between 6 and 8 by stability selection in all 8 cases and by the scaled minimax concave penalized selection for both $p = 200$ and $p = 3000$, these two methods

provide most consistent results. The scaled lasso and the adjusted cross-validation yield identical lasso and stability selections for $p = 200$ and identical stability selection for $p = 3000$.

We also compare the prediction performance of the scaled lasso with that of the lasso with the best fixed penalty level. We compute the scaled estimators in 1000 replications. In each replication, the dataset is split at random into a training set with 80 observations and a test set with 40 observations. The prediction mean squared error is computed within the test set, while the scaled estimators and the lasso estimator with fixed penalty level λ are computed based on the training set. Figure 1 demonstrates that in prediction, the scaled lasso with λ_0 chosen as $\lambda_2 = \{2(\log p)/n\}^{1/2}$ performs almost as well as the lasso with the optimal fixed λ .

In addition, we compare the prediction performance of all the estimators mentioned in this section. In each replication, we compute the penalized maximum likelihood estimator, its bias-correction, and scaled penalization methods based on the training set of 80 observations. For cross-validation, the training set of 80 observations is further partitioned at random 100 times into two groups of sizes 60 and 20, and a penalty level is selected by minimizing the estimated loss in the smaller group for the lasso estimator based on the larger group. This selected penalty level is then used for the lasso with the entire training set. Thus, the cross-validation lasso is also based on the training set with 80 observations. For the penalty level selected by the adjusted cross-validation, two estimators are considered: the lasso estimator and the least squares estimator after the lasso selection. In Table 4, we present the medians of the prediction mean squared error and the selected model size in the 200 replications. The scaled lasso has comparable prediction performance as cross-validation. Again, Table 4 suggests that original cross-validation tends to choose larger models, while adjusted cross-validation leads to results comparable with the scaled lasso.

5. DISCUSSION

In the theoretical analysis, we have considered $\lambda_0 = A\{(2/n)\log p\}^{1/2}$ with $A > 1$. This choice is somewhat conservative from a number of points of view. Simulation results suggest that the requirement $A > 1$ is a mathematical technicality. If $|X'\varepsilon/n|_\infty \leq \lambda_*$ with large probability for a standard normal vector ε , the theoretical results in this paper are all valid under $\text{pr}_{\beta, \sigma}$ when λ_0 is replaced by the smaller $\min(\lambda_0, A\lambda_*)$. The value of λ_* can be estimated by simulation with the given X and separately generated ε . A somewhat sharper theoretical choice of λ_0 is $A\{(2/n)\log(p/s)\}^{1/2}$ with the unknown $s = |\beta^*|_0$ (Zhang, 2010), or its simulated version with $\lambda_* = \max_{|T|=s} |X'_T \varepsilon|_2 / |T|^{1/2}$. The difference between the two λ_0 is limited unless $\log p = \{1 + o(1)\} \log n$. A reviewer called our attention to an unpublished 2011 report by Baraud, Giraud and Huet, available at <http://arxiv.org/abs/1007.2096>, whose method can be used to select a penalty level to nearly minimize the order of a penalized prediction error. This may also justify the use of smaller estimated penalty levels.

In the proof of our theoretical results for the scaled lasso, we use oracle inequalities for fixed penalty which unify and somewhat sharpen existing results. We now present this result. Define

$$\eta^*(\lambda, \xi) = \min_T 2^{-1} \left[\eta(\lambda, \xi, \beta^*, T) + \{ \eta^2(\lambda, \xi, \beta^*, T) - 16\lambda^2 |\beta_{T^c}^*|_2^2 \}^{1/2} \right] \quad (30)$$

as a sharper version of $\eta(\lambda, \xi, \beta^*, T)$ in (10).

THEOREM 4. *Let $\hat{\beta}(\lambda)$ be the minimizer of (1) with $\rho(t) = t$. Let $\beta^* \in \mathbb{R}^p$ be a target vector and $\xi > 1$. Then, in the event $|X'(y - X\beta^*)|_\infty/n \leq \lambda(\xi - 1)/(\xi + 1)$, we have*

$$|X\hat{\beta}(\lambda) - X\beta^*|_2^2/n \leq \min \{ \eta_*(\lambda, \xi), \eta^*(\lambda, \xi) \} \quad (31)$$

with $\eta_*(\lambda, \xi)$ in (12). Moreover, in the same event and with $\mu(\lambda, \xi)$ in (16),

$$|\hat{\beta}(\lambda) - \beta^*|_1 \leq \mu(\lambda, \xi). \quad (32)$$

The interpretations of (31) and (32) are given in (15) and (17), along with their relationship to several existing results. We note here that the condition $\kappa(\xi, S) \asymp 1$ for (15) and (17), weaker than the parallel condition on the restricted eigenvalue (Bickel et al., 2009), can be slightly weakened by using $F_1(\xi, S)$ in (21) (Ye & Zhang, 2010).

ACKNOWLEDGEMENT

This research was supported by the National Science Foundation and the National Security Agency. We thank Jian Huang for sharing the gene expression data, and reviewers for valuable suggestions.

APPENDIX

Here we prove Proposition 1, Theorem 4, Theorem 1, Theorem 2 and then Theorem 3.

Proof of Proposition 1. (i) Since $\hat{\beta} = \hat{\beta}(\sigma\lambda_0)$ is a solution of (2) at $\lambda = \sigma\lambda_0$,

$$\left\{ (\partial/\partial w)L_{\lambda_0}(w, \sigma) \Big|_{w=\hat{\beta}(\sigma\lambda_0)} \right\}_j = 0,$$

for all $\hat{\beta}_j(\sigma\lambda_0) \neq 0$. Since $\{j : \hat{\beta}_j(\lambda)\}$ is unchanged in a neighborhood of $\sigma\lambda_0$, $[(\partial/\partial\sigma)\{\hat{\beta}(\sigma\lambda_0)/\sigma\}]_j = 0$ for $\hat{\beta}_j(\sigma\lambda_0) = 0$. Thus,

$$\frac{\partial}{\partial\sigma}L_{\lambda_0}\{\hat{\beta}(\sigma\lambda_0), \sigma\} = \frac{\partial}{\partial t}L_{\lambda_0}\{\hat{\beta}(\sigma\lambda_0), t\} \Big|_{t=\sigma} = \frac{1-a}{2} - \frac{|y - X\hat{\beta}(\sigma\lambda_0)|_2^2}{2n\sigma^2}.$$

(ii) The convergence of (3) and (4) follows from the joint convexity of $L_{\lambda_0}(\beta, \sigma)$. The scale invariance follows from $L_0(c\beta, c\sigma; X, cy) = cL_0(\beta, \sigma; X, y)$, where $L_0(\beta, \sigma; X, y)$ expresses the dependence of (5) on the data (X, y) . \square

Proof of Theorem 4. (i) Let $\hat{\beta} = \hat{\beta}(\lambda)$. Since $\sigma^*z^* = |X'(y - X\beta^*)|_\infty/n$ and $\dot{\rho}(|\hat{\beta}_j|/\lambda) = 1$ for $\hat{\beta}_j \neq 0$, the inner product of $w - \hat{\beta}$ and the Karush–Kuhn–Tucker condition (2) yield

$$(X\hat{\beta} - Xw)'(X\hat{\beta} - X\beta^*)/n \leq \lambda(|w|_1 - |\hat{\beta}|_1) + \sigma^*z^*|w - \hat{\beta}|_1.$$

Since $2(X\hat{\beta} - Xw)'(X\hat{\beta} - X\beta^*) = |X\hat{\beta} - Xw|_2^2 + |Xh|_2^2 - |X\beta^* - Xw|_2^2$, this gives the basic inequality (25). Let $h = \hat{\beta} - \beta^*$. Since $\sigma^*z^* \leq \lambda(\xi - 1)/(\xi + 1)$, $\lambda\{|w|_1 - |\hat{\beta}|_1\} + \sigma^*z^*|w - \hat{\beta}|_1$ is no greater than $b|(w - \hat{\beta})_T|_1 + 2\lambda|w_{T^c}|_1 - (b/\xi)|(w - \hat{\beta})_{T^c}|_1$ with $b = 2\xi\lambda/(\xi + 1)$. Thus, (25) implies

$$|X\hat{\beta} - Xw|_2^2/n + |Xh|_2^2/n + (2b/\xi)|(w - \hat{\beta})_{T^c}|_1 \leq 2c + 2b|(w - \hat{\beta})_T|_1 \quad (A1)$$

with $c = |X\beta^* - Xw|_2^2/(2n) + 2\lambda|w_{T^c}|_1$. For $T = \emptyset$ and $w = \beta^*$, (A1) directly yields $|Xh|_2^2/n \leq c = 2\lambda|\beta^*|_1$. For general $\{w, T\}$, we want to prove

$$|Xh|_2^2/n \leq \eta(\lambda, \xi, w, T) = 2c + b^2/a, \quad a = \kappa^2(\xi, T)/|T|.$$

It suffices to consider $|Xh|_2^2/n \geq 2c$. In this case, $\hat{\beta} - w \in \mathcal{C}(\xi, T)$ by (A1), so that by (11)

$$a|(w - \hat{\beta})_T|_1^2 \leq |X\hat{\beta} - Xw|_2^2/n. \quad (A2)$$

Let $x = |(w - \hat{\beta})_T|_1$ and $y = |Xh|_2^2/n$. It follows from (A1) and (A2) that $ax^2 + y \leq 2c + 2bx$. For such (x, y) , $y - 2c \leq \max_x \{2bx - ax^2\} = b^2/a$. This gives $y \leq 2c + b^2/a = \eta(\lambda, \xi, w, T)$.

For $w = \beta^*$, it suffices to consider the case $y > c = 2\lambda|\beta_{T^c}^*|_1$, where the cone condition holds for $\hat{\beta} - \beta^*$. Now, (x, y) satisfies $ax^2 \leq y \leq c + bx$. The maximum of y , attained at $ax^2 = c + bx$, is

$$c + b\{b + (b^2 + 4ac)^{1/2}\}/(2a) = [\eta(\lambda, \xi, \beta^*, T) + \{\eta^2(\lambda, \xi, \beta^*, T) - 4c^2\}^{1/2}]/2. \quad (\text{A3})$$

(ii) Let $0 < \nu < 1$ and $T \subset \{1, \dots, p\}$. It follows from (A1) with $w = \beta^*$ that

$$(1 + \xi)|Xh|_2^2/n + 2\lambda|h_{T^c}|_1 \leq 2\lambda(\xi + 1)|\beta_{T^c}^*|_1 + 2\xi\lambda|h_T|_1.$$

It suffices to consider $\nu|h|_1 \geq (\xi + 1)|\beta_{T^c}^*|_1$. In this case

$$(1 + \xi)|Xh|_2^2/n + 2\lambda(1 - \nu)|h_{T^c}|_1 \leq 2\lambda(\xi + \nu)|h_T|_1.$$

Thus, $(1 - \nu)|h_{T^c}|_1 \leq (\xi + \nu)|h_T|_1$, or equivalently $h \in \mathcal{C}\{(\xi + \nu)/(1 - \nu), T\}$. It follows from (11) that $|Xh|_2^2/n \geq |h_T|_1^2 \kappa^2 \{(\xi + \nu)/(1 - \nu), T\}/|T|$, so that

$$(1 + \xi)|h_T|_1^2 \kappa^2 \{(\xi + \nu)/(1 - \nu), T\}/|T| + 2(1 - \nu)\lambda|h_{T^c}|_1 \leq 2(\xi + \nu)\lambda|h_T|_1. \quad (\text{A4})$$

Let $x = |h_T|_1$ and $y = |h_{T^c}|_1$. Write (A4) as $ax^2 + by \leq cx$. Subject to this inequality, the maximum of $x + y$ is $\max_{x \geq 0} \{x + (cx - ax^2)/b\}$. This maximum, attained at $2ax = b + c$, is $x(b + c)/(2b) = (b + c)^2/(4ab)$. Thus,

$$|h|_1 \leq \frac{\{2(\xi + 1)\lambda\}^2|T|}{4(1 + \xi)\kappa^2\{(\xi + \nu)/(1 - \nu), T\}\{2(1 - \nu)\lambda\}} = \frac{(\xi + 1)\lambda|T|/(1 - \nu)}{2\kappa^2\{(\xi + \nu)/(1 - \nu), T\}}.$$

This gives $|h|_1 \leq \mu(\lambda, \xi)$ for $\nu|h|_1 \geq (\xi + 1)|\beta_{T^c}^*|_1$. \square

Proof of Theorem 1. Assume $\tau_0 < 1$ without loss of generality. Consider $t \geq \sigma^*(1 - \tau_0)$ and the penalty level $\lambda = t\lambda_0$ for the lasso. Since $z^*\sigma^* \leq \sigma^*(1 - \tau_0)\lambda_0(\xi - 1)/(\xi + 1) \leq \lambda(\xi - 1)/(\xi + 1)$ and $\sigma^* = |y - X\beta^*|_2/n^{1/2}$, the Cauchy–Schwarz inequality and (31) imply

$$|y - X\hat{\beta}(t\lambda_0)|_2/n^{1/2} - \sigma^* \leq |X\hat{\beta}(t\lambda_0) - X\beta^*|_2/n^{1/2} \leq \eta_*^{1/2}(t\lambda_0, \xi).$$

Since $\eta_*^{1/2}(t\lambda_0, \xi) \leq \sigma^*\tau_0$ for $t < \sigma^*$, the derivative (7) of the loss with $a = 0$ satisfies

$$2t^2 \frac{\partial}{\partial t} L_{\lambda_0}\{\hat{\beta}(t\lambda_0), t\} = t^2 - |y - X\hat{\beta}(t\lambda_0)|_2^2/n \leq t^2 - (\sigma^*)^2(1 - \tau_0)^2 = 0 \text{ at } t = \sigma^*(1 - \tau_0).$$

This implies $\hat{\sigma} \geq \sigma^*(1 - \tau_0)$ by the strict convexity of the profile loss (5) in σ . For $t > \sigma^*$, $\eta_*^{1/2}(t\lambda_0, \xi) \leq t\tau_0$ by (10) and (12), so that at $t = \sigma^*/(1 - \tau_0)$,

$$t^2 - |y - X\hat{\beta}(t\lambda_0)|_2^2/n \geq t^2 - (\sigma^* + t\tau_0)^2 \geq 0.$$

This implies $\sigma^* \geq \hat{\sigma}(1 - \tau_0)$ by the strict convexity of (5) in σ . Thus, the first part of (13) holds. Moreover,

$$|X\hat{\beta} - X\beta^*|_2/n^{1/2} \leq \eta_*^{1/2}(\hat{\sigma}\lambda_0, \xi) \leq \eta_*^{1/2}\{\sigma^*\lambda_0/(1 - \tau_0), \xi\} \leq \sigma^*\tau_0/(1 - \tau_0).$$

Finally, since $\text{pr}_{\beta, \sigma}[|X'(y - X\beta)/n|_\infty \leq \sigma\{(2/n) \log p\}^{1/2}] \rightarrow 1$, (14) follows from (13). \square

The proof of Theorem 2 requires the following lemma.

LEMMA 1. Let T_m have the t -distribution with m degrees of freedom. Then, there exists $\epsilon_m \rightarrow 0$ such that for all $t > 0$

$$\text{pr}[T_m^2 > m\{e^{2t^2/(m-1)} - 1\}] \leq (1 + \epsilon_m)e^{-t^2}/(\pi^{1/2}t). \quad (\text{A5})$$

Proof of Lemma 1. Let $x = [m\{e^{2t^2/(m-1)} - 1\}]^{1/2}$. Since T_m has the t -distribution,

$$\begin{aligned} \text{pr}(T_m^2 > x^2) &= \frac{2\Gamma\{(m+1)/2\}}{\Gamma(m/2)(m\pi)^{1/2}} \int_x^\infty \left(1 + \frac{u^2}{m}\right)^{-(m+1)/2} du \\ &\leq \frac{2\Gamma\{(m+1)/2\}}{x\Gamma(m/2)(m\pi)^{1/2}} \int_x^\infty \left(1 + \frac{u^2}{m}\right)^{-(m+1)/2} u du \\ &= \frac{2\Gamma\{(m+1)/2\}m}{x\Gamma(m/2)(m\pi)^{1/2}(m-1)} \left(1 + \frac{x^2}{m}\right)^{-(m-1)/2}. \end{aligned}$$

Since $x \geq t\{2m/(m-1)\}^{1/2}$,

$$\text{pr}(T_m^2 > x^2) \leq \frac{\sqrt{2}\Gamma\{(m+1)/2\}}{\Gamma(m/2)(m-1)^{1/2}} \frac{e^{-t^2}}{t\pi^{1/2}} = (1 + \epsilon_m) \frac{e^{-t^2}}{t\pi^{1/2}},$$

where $\epsilon_m = \{2/(m-1)\}^{1/2}\Gamma\{(m+1)/2\}/\Gamma(m/2) - 1 \rightarrow 0$ as $m \rightarrow \infty$. \square

Proof of Theorem 2. We need to express τ_*^2 as a function of σ at $\sigma = \sigma^*$ in the proof. Define

$$\phi(\sigma) = \lambda_0 \mu(\sigma \lambda_0, \xi) / \sigma, \quad \phi_+ = \frac{\phi(\sigma^*)\xi}{(\xi+1)\{1-\phi(\sigma^*)\}_+}, \quad \phi_- = \frac{\phi(\sigma^*)(\xi-1)}{\xi+1}.$$

We have $\tau_*^2 = \phi(\sigma^*) < 1$, $\phi_- \leq \phi(\sigma^*)$ and $\phi_+ \leq \phi(\sigma^*)/(1-\phi(\sigma^*))$.

(i) Consider $z^* \leq (1-\phi_-)\lambda_0(\xi-1)/(\xi+1)$. Let $\lambda = t\lambda_0$ and $h = \hat{\beta}(\lambda) - \beta^*$. Since $|X'(y - X\beta^*)/n|_\infty = z^*\sigma^*$, the Karush–Kuhn–Tucker condition (2) gives

$$\begin{aligned} -(z^*\sigma^* + \lambda)|h|_1 &\leq (Xh)' \{y - X\beta^* + y - X\hat{\beta}(\lambda)\} / n \\ &= (\sigma^*)^2 - |y - X\hat{\beta}(\lambda)|_2^2 / n \\ &= (Xh)' \{2(y - X\beta^*) - Xh\} / n \leq 2z^*\sigma^*|h|_1 \end{aligned} \quad (\text{A6})$$

as lower and upper bounds for $(\sigma^*)^2 - |y - X\hat{\beta}(\lambda)|_2^2 / n$. This is a key point in the proof.

For $t \geq \sigma^*(1-\phi_-)$, $z^*\sigma^* \leq t\lambda_0(\xi-1)/(\xi+1) = \lambda(\xi-1)/(\xi+1)$, so that (32) in Theorem 4 implies $|h|_1 \leq \mu(t\lambda_0, \xi)$. It follows (A6) that for $t = \sigma^*(1-\phi_-)$,

$$\begin{aligned} t^2 - |y - X\hat{\beta}(t\lambda_0)|_2^2 / n &\leq t^2 - (\sigma^*)^2 + 2z^*\sigma^*\mu(t\lambda_0, \xi) \\ &\leq 2t(t - \sigma^*) + 2t\lambda_0(\xi-1)(\xi+1)^{-1}\mu(\sigma^*\lambda_0, \xi) = 0, \end{aligned}$$

due to $\phi_- = (\xi-1)(\xi+1)^{-1}\phi(\sigma^*) = (\xi-1)(\xi+1)^{-1}\lambda_0\mu(\sigma^*\lambda_0, \xi)/\sigma^*$. As in the proof of Theorem 1, we find $\hat{\sigma}/\sigma^* \geq 1 - \phi_-$ by (7) and the strict convexity of (5) in σ .

Now we prove that $\hat{\sigma}/\sigma^* \leq 1 + \phi_+$. For $t > \sigma^*$, $\mu(t\lambda_0, \xi) \leq (t/\sigma^*)\mu(\sigma^*\lambda_0, \xi)$ by (16). Thus, since $(\xi-1)/(\xi+1) + 1 = 2\phi_+\{1-\phi(\sigma^*)\}/\phi(\sigma^*)$ and $\phi_+ \leq (1+\phi_+)\phi(\sigma^*)$, for $t/\sigma^* = 1 + \phi_+$, (A6) and (32) imply that

$$\begin{aligned} t^2 - |y - X\hat{\beta}(t\lambda_0)|_2^2 / n &\geq t^2 - (\sigma^*)^2 - (z^*\sigma^* + t\lambda_0)\mu(t\lambda_0, \xi) \\ &\geq (t + \sigma^*)\sigma^*\phi_+ - \{(\xi-1)/(\xi+1) + 1 + \phi_+\}t\lambda_0\mu(\sigma^*\lambda_0, \xi) \\ &= (\sigma^*)^2\{(2 + \phi_+)\phi_+ - [2\phi_+\{1-\phi(\sigma^*)\}/\phi(\sigma^*) + \phi_+](1 + \phi_+)\phi(\sigma^*)\} \\ &= (\sigma^*)^2\phi_+\{\phi(\sigma^*)(1 + \phi_+) - \phi_+\} > 0. \end{aligned}$$

It follows that $\hat{\sigma}/\sigma^* \leq 1 + \phi_+$ by convexity.

Since $1 - \phi_- \leq \hat{\sigma}/\sigma^* \leq 1 + \phi_+$, $|\hat{\beta}(\hat{\sigma}\lambda_0) - \beta^*|_1 \leq \mu(\hat{\sigma}\lambda_0, \xi) \leq \mu(\sigma^*\lambda_0, \xi)(1 + \phi_+)$. This completes the proof of (18).

(ii) Let $z_j = x'_j(y - X\beta^*)/(n\sigma^*)$ with $z^* = \max_{j \leq p} |z_j|$. Under $\text{pr}_{\beta^*, \sigma^*}$, $\varepsilon^* = y - X\beta^*$ is a vector of independent and identically distributed normal variables with zero mean. Since $\sigma^* = |y - X\beta^*|/n^{1/2}$, $z_j/\{(1 - z_j^2)/(n-1)\}^{1/2}$ follows a t -distribution with $n-1$ degrees of freedom. Lemma 1 with $m =$

$n - 1$ and $t^2 = \log(p/\epsilon) > 2$ implies

$$\text{pr}_{\beta^*, \sigma} \left[\frac{(n-1)z_j^2}{1-z_j^2} > (n-1)\{e^{2t^2/(n-2)} - 1\} \right] \leq \frac{1 + \epsilon_{n-1}}{\pi^{1/2}t} e^{-t^2} = \frac{(1 + \epsilon_{n-1})\epsilon/p}{\{\pi \log(p/\epsilon)\}^{1/2}}. \quad (\text{A7})$$

Since $e^a - 1 \leq \sum_{k=1}^{\infty} a^k / 2^{k-1} = a/(1 - a/2)$ for any $0 < a < 2$,

$$(n-1)\{e^{2t^2/(n-2)} - 1\} \leq \frac{2(n-1)t^2/(n-2)}{1 - t^2/(n-2)} \leq \frac{2(n-1)t^2/n}{1 - 2t^2/n}. \quad (\text{A8})$$

The combination of (A7) and (A8) yields

$$\begin{aligned} \text{pr}_{\beta^*, \sigma} [|z_j| > \{2 \log(p/\epsilon)/n\}^{1/2}] &= \text{pr}_{\beta^*, \sigma} \left\{ \frac{(n-1)z_j^2}{1-z_j^2} > \frac{2(n-1)t^2/n}{1-2t^2/n} \right\} \\ &\leq \text{pr}_{\beta^*, \sigma} \left\{ \frac{(n-1)z_j^2}{1-z_j^2} > (n-1)(e^{\frac{2t^2}{n-2}} - 1) \right\} \\ &\leq (1 + \epsilon_{n-1})(\epsilon/p)/\{\pi \log(p/\epsilon)\}^{1/2}. \end{aligned}$$

Since $\lambda_0 \geq \{(2/n) \log(p/\epsilon)\}^{1/2}(\xi + 1)/\{(\xi - 1)(1 - \phi_-)\}$, this bounds the tail probability of $z^* = \max_{j \leq p} |z_j|$ by the union bound. Since $n(\sigma^*/\sigma)^2$ follows the χ_n^2 distribution, $n^{1/2}(\sigma^*/\sigma - 1)$ converges to $N(0, 1/2)$ in distribution, which then implies (19) by (18) under $\phi(\sigma) = o(n^{-1/2})$. \square

Proof of Theorem 3. Let $h = \hat{\beta} - \beta^*$. It follows from the proof of Theorem 2 (i) that $z^* \leq (1 - \phi_-)\lambda_0(\xi - 1)/(\xi + 1) \leq (\hat{\sigma}/\sigma^*)\lambda_0(\xi - 1)/(\xi + 1)$, so that $\hat{\lambda} + z^*\sigma^* \leq \xi(\hat{\lambda} - z^*\sigma^*)$. By (2), $|x'_j Xh/n| = |x'_j(y - X\hat{\beta} - \varepsilon^*)/n| \geq \hat{\lambda} - z^*\sigma^*$ for $\hat{\beta}_j \neq 0$. Let $B \subseteq \hat{S} \setminus S$ with $|B| \leq m$. Since $\kappa_+(m, S)$ is the upper sparse eigenvalue, $(\hat{\lambda} - z^*\sigma^*)^2|B| \leq \kappa_+(m, S)|Xh|_2^2/n$. By the basic inequality (25) with $w = \beta^*$, $|Xh|_2^2/n \leq (\hat{\lambda} + z^*\sigma^*)|h_S|_1$ and h is in the cone $\mathcal{C}(\xi, S)$. Thus, since $|h_S|_1^2 \kappa^2(\xi, S) \leq |Xh|_2^2|S|/n$ by (11), $|Xh|_2^2/n \leq (\hat{\lambda} + z^*\sigma^*)^2|S|/\kappa^2(\xi, S)$. It follows that $|B| \leq \kappa_+(m, S)\xi^2|S|/\kappa^2(\xi, S) < m$. Since all $B \subseteq \hat{S} \setminus S$ of size $|B| \leq m$ have size $|B| < m$, $\hat{S} \setminus S$ does not have a subset of size m . This gives the first inequality in (27).

Let P_B be the orthogonal projection to the linear span of $(x_j, j \in B)$. By the definition of $\sigma_{m,S}^*$ and the prediction error bound $|Xh|_2^2/n \leq \eta_*(\hat{\lambda}, \xi)$ in Theorem 4,

$$\hat{\sigma}^2 - \bar{\sigma}^2 = |P_{\hat{S}}(y - X\hat{\beta})|_2^2/n \leq (|P_{\hat{S}}\varepsilon^*|_2 + |P_{\hat{S}}Xh|_2)^2/n \leq \{\sigma_{m-1,S}^* + \sqrt{\eta_*(\hat{\lambda}, \xi)}\}^2.$$

This gives the second inequality in (27). The prediction error bound follows from

$$|X\bar{\beta} - X\beta^*|_2 = |P_{\hat{S}}y - X\beta^*|_2 \leq |P_{\hat{S}}(y - X\hat{\beta})|_2 + |Xh|_2 \leq \{\sigma_{m-1,S}^* + 2\sqrt{\eta_*(\hat{\lambda}, \xi)}\}\sqrt{n},$$

which implies the ℓ_2 estimation error bound due to $\kappa_-(m-1, S)|\bar{\beta} - \beta^*|_2^2 \leq |X\bar{\beta} - X\beta^*|_2^2/n$. Finally, the probability bound in (29) follows directly from an application of the Gaussian concentration inequality to the chi-squared variables in the union bound. \square

REFERENCES

- ANTONIADIS, A. (2010). Comments on: ℓ_1 -penalization for mixture regression models by N. Städler, P. Bühlmann and S. van de Geer. *Test* **19**, 257–258.
- BELLONI, A., CHERNOZHUKOV, V. & WANG, L. (2011). Square-root lasso: Pivotal recovery of sparse signals via conic programming. *Biometrika* **98**, 791–806.
- BICKEL, P., RITOV, Y. & TSYBAKOV, A. (2009). Simultaneous analysis of lasso and Dantzig selector. *Annals of Statistics* **37**, 1705–1732.
- BUNEA, F., TSYBAKOV, A. & WEGKAMP, M. H. (2007). Sparsity oracle inequalities for the lasso. *Electronic Journal of Statistics* **1**, 169–194.
- CANDES, E. & TAO, T. (2007). The Dantzig selector: statistical estimation when p is much larger than n (with discussion). *Annals of Statistics* **35**, 2313–2404.

- EFRON, B., HASTIE, T., JOHNSTONE, I. & TIBSHIRANI, R. (2004). Least angle regression (with discussion). *Annals of Statistics* **32**, 407–499.
- FAN, J., GUO, S. & HAO, N. (2012). Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *Journal of Royal Statistical Society B* **74**, 37–65.
- FAN, J. & LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* **96**, 1348–1360.
- FAN, J. & PENG, H. (2004). On non-concave penalized likelihood with diverging number of parameters. *Annals of Statistics* **32**, 928–961.
- GREENSHTEIN, E. (2006). Best subset selection, persistence in high-dimensional statistical learning and optimization under ℓ_1 constraint. *Annals of Statistics* **34**, 2367–2386.
- GREENSHTEIN, E. & RITOV, Y. (2004). Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. *Bernoulli* **10**, 971–988.
- HUANG, J., MA, S. & ZHANG, C.-H. (2008). Adaptive lasso for sparse high-dimensional regression models. *Statistica Sinica* **18**, 1603–1618.
- HUBER, P. J. & RONCHETTI, E. M. (2009). *Robust Statistics*. Wiley, 2nd ed., pp. 172–175.
- KOLTCHINSKII, V., LOUNICI, K. & TSYBAKOV, A. B. (2011). Nuclear norm penalization and optimal rates for noisy low rank matrix completion. *Annals of Statistics* **39**, 2302–2329.
- MEINSHAUSEN, N. & BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *Annals of Statistics* **34**, 1436–1462.
- MEINSHAUSEN, N. & BÜHLMANN, P. (2010). Stability selection (with discussion). *Journal of the Royal Statistical Society: Series B* **72**, 417–473.
- MEINSHAUSEN, N. & YU, B. (2009). Lasso-type recovery of sparse representations for high-dimensional data. *Annals of Statistics* **37**, 246–270.
- OSBORNE, M., PRESNELL, B. & TURLACH, B. (2000a). A new approach to variable selection in least squares problems. *IMA Journal of Numerical Analysis* **20**, 389–404.
- OSBORNE, M., PRESNELL, B. & TURLACH, B. (2000b). On the lasso and its dual. *Journal of Computational and Graphical Statistics* **9**, 319–337.
- SCHEETZ, T. E., KIM, K.-Y. A., SWIDERSKI, R. E., PHILP, A. R., BRAUN, T. A., KNUDTSON, K. L., DORRANCE, A. M., DiBONA, G. F., HUANG, J., CASAVANT, T. L., SHEFFIELD, V. C. & STONE, E. M. (2006). Regulation of gene expression in the mammalian eye and its relevance to eye disease. *Proc. Nat. Acad. Sci* **103**, 14429–14434.
- SHE, Y. & OWEN, A. B. (2011). Outlier detection using nonconvex penalized regression. *Journal of the American Statistical Association* **106**, 626–639.
- STÄDLER, N., BÜHLMANN, P. & VAN DE GEER, S. (2010). ℓ_1 -penalization for mixture regression models (with discussion). *Test* **19**, 209–285.
- SUN, T. & ZHANG, C.-H. (2010). Comments on: ℓ_1 -penalization for mixture regression models by N. Städler, P. Bühlmann and S. van de Geer. *Test* **19**, 270–275.
- VAN DE GEER, S. (2008). High-dimensional generalized linear models and the lasso. *Annals of Statistics* **36**, 614–645.
- VAN DE GEER, S. & BÜHLMANN, P. (2009). On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics* **3**, 1360–1392.
- YE, F. & ZHANG, C.-H. (2010). Rate minimaxity of the lasso and Dantzig selector for the ℓ_q loss in ℓ_r balls. *Journal of Machine Learning Research* **11**, 3481–3502.
- ZHANG, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics* **38**, 894–942.
- ZHANG, C.-H. & HUANG, J. (2008). The sparsity and bias of the lasso selection in high-dimensional linear regression. *Annals of Statistics* **36**, 1567–1594.
- ZHANG, T. (2009). Some sharp performance bounds for least squares regression with L_1 regularization. *Ann. Statist.* **37**, 2109–2144.
- ZHAO, P. & YU, B. (2006). On model selection consistency of lasso. *Journal of Machine Learning Research* **7**, 2541–2567.

[Received mmm yyyy. Revised mmm yyyy]